arXiv:1402.1054v1 [stat.AP] 5 Feb 2014

ON RECENT ADVANCES IN SUPERVISED RANKING FOR METABOLITE PROFILING

CHARANPAL DHANJAL AND STÉPHAN CLÉMENÇON

Institut Mines-Télécom; Télécom ParisTech, CNRS LTCI, 46 rue Barrault, 75634 Paris Cedex 13, France

Abstract. This paper focuses on data arising from the field of *metabolomics*, a rapidly developing area concerned by the analysis of the chemical fingerprints (*i.e.* the metabolite profile). The metabolite profile is left by specific chemical processes occurring in biological cells, tissues or organs. It is the main purpose of this article to develop and implement scoring techniques so as to rank all possible *metabolic profiles* by increasing order of magnitude of the conditional probability that a given metabolite is present at high levels in a certain biological fluid. After a detailed description of the (functional) data from which decision rules must be learnt, several approaches to this predictive problem, based on recent advances in K-partite ranking are described at length. Their performance on several real datasets are next thoroughly investigated.

metabolite profiling, functional data analysis, regularisation, filtering, bi-partite ranking

1. Introduction

Thanks to high resolution *Mass Spectrometry* (MS) and *Nuclear Magnetic Resonance* (NMR), it is now possible to get very detailed metabolic information about the chemical processes which might occur in a biological cell, tissue, organ or organism. *Metabolomics* is precisely the study of these biological processes through the quantitative measurements obtained by means of one of these analytical platforms. It belongs to the most recent -omics techniques used in order to *phenotype* microorganisms or multicellular organisms like plants or animals, *i.e.* to characterise the metabolic response of such living systems to given pathophysiological stimuli or genetic modification, see [33, 16, 36] for instance.

Whereas MS measures the relative abundance in ionised analytes, directly or after a chromatographic step, which prevents some signal extinction problems and warrants a better quantification of the few ionised species simultaneously detected at a given retention time, NMR is less sensitive but more rapid to implement and needs a very simple sample preparation (the term "metabonomics" is sometimes used to distinguish analyses based on NMR measurements). NMR detection of analytes corresponds to the detection of the spin of non-exchangeable protons, or carbon-13, present on analytes. Due to the presence of chemical neighbouring of specific chemical functions, assignments of these spins is specific for a given analyte.

E-mail address: {charanpal.dhanjal, stephan.clemencon}@telecom-paristech.fr.
Date: February 6, 2014.

Nevertheless, in biological matrices, due to the coalescence of some signals coming from numerous analytes in certain spectral regions, the ultimate identification can be difficult to obtain in one dimensional NMR. For any individual submitted to this phenotyping procedure, this information is summarised inside a spectroscopic fingerprint, which contains hundreds to thousands of analytes. They are a priori considered as independent variables. Mining such a large information requires sophisticated statistical tools [38, 15, 14]. This represents a highly challenging domain in statistics research, most tools being still at a conceptual or a developing stage.

Contrary to MS-based fingerprints, NMR ones can be obtained easily on very large populations with a high reproducibility and therefore are well adapted to an epidemiological survey of the general population or of some targeted groups, which display some pathological risks [26]. Indeed, this is the case of high level or elite sportsmen who are submitted in France since beginning of 2000s to the longitudinal medical follow-up to detect endocrine anomalies which could be indicative of forbidden doping practices. In bike athletes, a biobank containing sera of sportsmen submitted to this longitudinal follow-up has been progressively built. In this biobank are present some sera in which extreme (lower or higher than normal) values of circulating cortisol, Insulin-like growth factor 1 (IGF-1), or testosterone are well documented. To get sufficient numbers of individuals inside the normal and abnormal classes for any of the 3 hormones studied, a biobank containing 655 samples was built from a cohort of 250 sportsmen.

Existing research on automatic prediction within metabolomic data has studied the use of classification algorithms such as Support Vector Machines (SVMs), Random Forests and Partial Least Squares (PLS, [39]), and subspace projection methodologies such as Principal Components Analysis (PCA, [27]). In our study we motivate the problem under the framework of bipartite ranking (explained in detail in the next section), and examine the data described above using a selection of state-of-the-art algorithms using a plethora of feature transformations which together account for the functional nature of the data.

The paper is organised as follows. In Section 2 we detail some of the challenges in making functional prediction and then introduce bipartite ranking as an informative way to view the prediction problem. Section 3 details our data corpus and the following section presents an empirical examination of the dataset. We conclude in Section 5.

2. Metabolomics Meets Machine-Learning

It is the purpose of this section to formulate the predictive problems related to metabolomics data from a machine-learning perspective and describe the nature of the input measurements, together with the difficulties raised by their (functional) nature.

2.1. NMR-based Fingerprints and Predictive Issues.

NMR-based fingerprints are recorded as a continuous spectrum however they are often transformed into a set of features using a bucketing technique. One of the challenges of working with this type of data is that it is high dimensional and the features are correlated and hence learning suffers from the curve of dimensionality (see [1] for a discussion). The classification based on curves such as those obtained using NMR is called *functional classification* since the observations are sampled functions rather than simple

high-dimensional vectors. The analysis of functional data is classed as Functional Data Analysis [34, 17].

One way to deal with the challenges faced with functional data is to use *dimensionality reduction* in which the inputs curves are projected into a finite dimensional subspace. Common examples of dimensionality reduction methods include PCA and PLS. The projection can then be used in conjunction with a classification algorithm to make predictions. Another commonly used approach in signal processing is *filtering*, which thresholds the coefficients of the spectrum that are below a certain base-line. As a simple example, one chooses the first d coefficients with respect to a certain basis. Yet another way to deal with functional data is to use *regularisation*, for example see [23].

2.2. **Wavelets.** As we mentioned, the direct use of statistical methods on raw spectra can result in poor performance, and to address this issue the *wavelet transform* [30, 10] can be used. Wavelets map spectra onto a set of frequency components with localisation in space as well as frequency. They can be used to compactly represent a signal, or via thresholding, to denoise signals. An advantage of wavelets over the another similar transform, the Fourier transform [3], is that wavelet coefficients are localised in time and hence frequency changes in the spectrum, such as discontinuities and peaks, can be effectively modelled. In contrast, Fourier analysis is conducted over the whole time period of the spectrum.

Consider the space of square integrable functions on $[0,1]$ denoted by $L_2([0,1])$, and denote the observed value of a variable Z at time t using the notation $Z(t)$. The wavelet decomposition finds coefficients at different times and scales to approximate L_2 with respect to a set of orthonormal functions $\psi_{j,k}$, $k = 0, \ldots, 2^j - 1$, using a *mother wavelet* $\psi_{j,k}(t) = 2^{j/2}\psi(2^j t - k)$ for resolution level or scale $j \geq 0$ and translation $k = 0, \ldots, 2^j - 1$. The space $L_2([0,1])$ can be approximated using a set of closed subspaces $V_0 \subset V_1 \cdots \subset L_2([0,1])$ in which V_j is spanned by 2^j orthonormal scaling functions $\phi_{j,k}$, $k = 0, \ldots, 2^j - 1$. The orthonormal complement between V_j and V_{j+1} is generated by the mother wavelets $\psi_{j,k}, k = 0, \ldots, 2^j - 1$. It follows that the space $L_2([0,1])$ can be written using the orthonormal basis $\cup_{j\geq 0}\{\psi_{j,k}\}_{k=0,\ldots,2^j-1}$ and $\phi_{0,0}$. Therefore an observation in Z can be written as

$$Z(t) = \sum_{j=0}^{\infty}\sum_{k=0}^{2^j-1} \alpha_{j,k}\psi_{j,k}(t) + \mu\phi_{0,0}(t),$$

for $t \in [0,1]$. The coefficients are found using $\alpha_{j,k} = \int_0^1 Z(t)\psi_{j,k}(t)dt$ and $\mu = \int_0^1 Z(t)\phi_{0,0}(t)dt$, which amounts to the projection of a spectra onto the corresponding basis function. In practice one limits the resolution level with an upper bound J and the approximation of the spectrum is given by

$$Z(t) \approx \sum_{j=0}^{J-1}\sum_{k=0}^{2^j-1} \alpha_{j,k}\psi_{j,k}(t) + \mu\phi_{0,0}(t).$$

Thus, instead of using the raw spectra features, one can use wavelet coefficients $\alpha_{j,k}$, $j = 0, \ldots, J-1$, $k = 0, \ldots, 2^j - 1$, and μ. Common mother wavelets include Haar, Daubechies and Symmlet wavelets. The latter two have different forms according to the number of *vanishing moments* in which a wavelet with M vanishing moments

has $\int t^k \psi(t) dt = 0$ for $k \in \{0, \ldots, M-1\}$. Wavelets have been used to preprocess metabolomic data in [11, 40].

2.3. **Prediction Approaches.** One way of tackling the classification problem using wavelets is to threshold the wavelet coefficients that have the smallest values, which are deemed irrelevant. In hard-thresholding one removes the smallest coefficients and leaves the others unchanged. Soft thresholding in addition reduces the values of the remaining coefficients by the threshold value. A more pragmatic approach is to find those coefficients which are most useful for prediction, for example in [2]. This effectively corresponds to *feature selection* [22], which is the problem of selecting a set of features which are relevant to the label being predicted.

Feature selection is a broad field ranging from individual feature scoring, to subset selection through search, to projection methods into subspaces. We consider combined feature selection and prediction approaches, and a useful technique is to find the minimum distance projection onto a subspace with a penalty on the L0 or L1 norm. Doing so results in a sparse set of projection vectors and hence not all features are used in the prediction of new test examples. In this vein, there have been a plethora of dimensionality reduction methods, for example Sparse PCA [42, 9, 31] and sparse PLS [25, 12, 32].

The Least Absolute Shrinkage and Selection Operator (LASSO [37]) is one approach for prediction that has gained popularity in recent years. When working with classification problems one can simply use the predicted real output and then threshold in order to provide a positive or negative label. LASSO is applied to metabolic data in [35] and compared to bootstrapped LASSO. Other prediction studies of metabolomic data include [13] which applies PLS Discriminant Analysis (PLS-DA) to MRI associated with orphan neurological disease. A study of of NMR spectra of urine of Streptococcus Pneumoniae in [29] demonstrates an improvement of SVMs over PLS-DA. In [5] Random Forests are applied to the metabolic classification problem, with a particular relevance to the important features used in prediction. In contrast to these studies, we examine metabolomic data in the framework of bipartite ranking which has the advantages outlined below.

2.4. **Bipartite Ranking.** The metabolomic data studied in this article is labelled with unbalanced classes, for example few positive and many negative labels. In this case it is advantageous to study the bipartite ranking problem as opposed to binary classification. In bipartite ranking one assigns scores to a set of examples labelled in a binary fashion. This is useful, as will later become clearer, since it monitors two error types: the rate of missed detections and false positives. The competitiveness of ranking compared to regression methods such as LogitBoost [20] and Kernel Logistic Regression [41] is demonstrated in [6].

To formalise bipartite ranking consider a set of examples-label pairs denoted by $S = \{(\mathbf{x}_1, y_1), \ldots, (\mathbf{x}_n, y_n)\}$ in which examples are denoted $\mathbf{x}_i \in \mathcal{X}$ and labels using $y_i \in \{-1, +1\}$, $i = 1, \ldots, n$. There exists a scoring function on the set of examples $s : \mathcal{X} \to \mathbb{R}$ such that $\forall (\mathbf{x}, \mathbf{x}') \in \mathcal{X}^2, \mathbf{x} \leq_s \mathbf{x}' \Leftrightarrow s(\mathbf{x}) \leq s(\mathbf{x}')$. The Receiver Operator Characteristic (ROC) curve $\mathrm{ROC}(s, t)$ is defined as $t \in \mathbb{R} \mapsto (P\{s(X) > t \mid Y = -1\}, P\{s(X) > t \mid Y = +1\})$ where t is a threshold on the scoring function. In words, the ROC curve is a plot of the true positive rate (probability of a positive prediction given a positive label) versus the false positive rate (probability of a positive prediction given a negative label) as the threshold value t changes. The

ROC curve is informative because it represents the classification errors for all scoring thresholds t. Hence, one can for example, choose an acceptable false positive rate and trade off the true positive rate accordingly.

One way to measure the quality of a ROC curve in a scalar manner, which is useful for optimisation, is using the Area Under the ROC Curve (AUC), defined by $\mathrm{AUC}(s) = \int_0^1 \mathrm{ROC}(s,\alpha)d\alpha$. A larger AUC corresponds to a better scoring function. It can also be written in terms of the rate of concording pairs,

$$\begin{aligned}\mathrm{AUC}(s) &= P(s(X) > s(X')|Y = +1, Y' = -1) \\ &+ \frac{1}{2}P(s(X) = s(X')|Y = +1, Y' = -1),\end{aligned}$$

where (X, Y) and (X', Y') are i.i.d. observations. In bipartite ranking the aim is to find a scoring function with a ROC curve as close to the optimal curve as possible and this implies a small residual AUC with respect to the optimal. In the following subsections we outline the bipartite ranking algorithms: Ranking Forests [8], Rank SVM [28] and RankBoost [18].

2.4.1. *Ranking Forests.* *TreeRank Forest* or *Ranking Forests* is a ranking algorithm based on a collection of individual *TreeRank* learners. One first constructs a set of *ranking trees* each denoted $\mathcal{T}_D$ where $D \geq 0$ is the depth of the tree. The root of the tree is labelled $C_{0,0}$ and the non-leaf nodes are given by $C_{d,k}$ for depth d and $0 \leq k \leq 2^d$, which has two children $C_{d+1,2k}$ (left) and $C_{d+1,2k+1} = C_{d,k}\backslash C_{d+1,2k}$ (right). Given a particular ranking tree the score of an example is computed as

$$s_{\mathcal{T}}(\mathbf{x}) = \sum_{C_{d,k}:\text{terminal cell}} 2^D(1 - k/2^d) \cdot \mathbb{I}\{\mathbf{x} \in C_{d,k}\}.$$

In other words, the scores ascend going from left to right along the leaf nodes.

To construct the tree, TreeRank follows a similar idea to that of Classification and Regression Trees (CART, [4]) and splits the data in order to maximise the AUC at each node. At cell C one split the data into two mutually exclusive subsets C_+ and $C_- = C\backslash C_+$ so that the error given by:

$$\begin{aligned}L_{C,\omega}(\Gamma) &= \frac{2(1-\omega)}{n}\sum_{i=1}^{n}\mathbb{I}\{\mathbf{x}_i \in C\backslash\Gamma\} \cdot \mathbb{I}\{y_i = +1\} \\ &+ \frac{2\omega}{n}\sum_{i=1}^{n}\mathbb{I}\{\mathbf{x}_i \in \Gamma \cap C\} \cdot \mathbb{I}\{y_i = -1\},\end{aligned}$$

(or a penalised/convexified variant) is minimised for a given weight $0 \leq \omega \leq 1$. One starts with $C_{0,0} = \mathcal{X}$ and the iterates over $d = 0, \ldots, D-1$ and $k = 0, \ldots, 2^d - 1$ in the following way: the rate of positives in the current cell $C_{d,k}$ is given by $\alpha(C_{d,k}) = \frac{1}{n_{d,k}}\sum_{i=1}^{n}\mathbb{I}\{\mathbf{x}_i \in C_{d,k}\} \cdot \mathbb{I}\{y_i = +1\}$, where $n_{d,k}$ is the number of elements in the cell. One then minimises $L_{C_{d,k},\omega}(\Gamma)$ with $\omega = \alpha(C_{d,k})$. This provides the split into the two child cells as $C_{d+1,2k} = C_+$ and $C_{d+1,2k+1} = C_{d,k}\backslash C_{d+1,2k}$. The splitting rule can be formed using any cost-sensitive classifier (see [24]) and termed a LeafRank algorithm. The algorithm is shown to be a piecewise linear interpolation scheme of the optimal ROC curve.

To combine ranking trees one uses bootstrap aggregating (bagging) in conjunction with feature randomisation. In essence, one draws a number B of bootstrap samples (random sampling with replacement) and then selects a random subset of

the features of the examples for each sample and these are used to train B ranking trees. There are a number of methods to aggregate the scoring functions of the set of trees, for example by computing the median ranking with respect to a distance metric between rankings. One can also simply compute the mean scoring function.

2.4.2. *RankBoost.* RankBoost is based on AdaBoost [19], and works by combining many "weak" rankings into a single accurate one. At iteration t one maintains a distribution matrix $\mathbf{D}_t \in \mathcal{X} \times \mathcal{X}$ which emphasises the importance of orderings on particular pairs of examples. We replicate the pseudo-code for RankBoost in Algorithm 1 for clarity.

Algorithm 1 RankBoost pseudo-code

1: Input: Initial distribution matrix $\mathbf{D}$
2: $\mathbf{D}_0 = \mathbf{D}$
3: **for** $t = 1, \ldots, T$ **do**
4: Train weak learner using $\mathbf{D}_t$ and choose $\alpha_t \in \mathbb{R}$
5: Get weak ranking $h_t : \mathcal{X} \to \mathbb{R}$
6: Set $\mathbf{D}_{t+1}(\mathbf{x}_0, \mathbf{x}_1) = \frac{\mathbf{D}_t(\mathbf{x}_0,\mathbf{x}_1)\exp(\alpha_t(h_t(\mathbf{x}_0)-h_t(\mathbf{x}_1)))}{\mathbf{Z}_t}$ where $\mathbf{Z}_t$ is a normalisation constant
7: **end for**
8: Output: $H(\mathbf{x}) = \sum_{i=1}^{T} \alpha_i h_i(\mathbf{x})$

The key idea is that the distribution matrix $\mathbf{D}_t$ is modified at each iteration so as to stress the importance of certain pair orderings. If $\mathbf{x}_1$ is to be ranked higher than $\mathbf{x}_0$ and $\alpha_t > 0$ then if $h_t(\mathbf{x}_1) > h_t(\mathbf{x}_0)$ the update rule on $\mathbf{D}_{t+1}(\mathbf{x}_0, \mathbf{x}_1)$ will increase its value. Further details of how to set α_t and some algorithms for weak learners are given in [18].

2.4.3. *Ranking SVM.* Finally we overview the Ranking SVM which is a way to maximise a functional closely related to an similarity measure between two rankings known as the *Kendall tau*. Let P be the number of concording pairs and Q be the number of discording pairs between two rankings h and h', then Kendall tau is defined as $\tau(h, h') = (P - Q)/(P + Q)$. The Ranking SVM approach is based on the idea of maximising a margin as in SVMs whilst minimising Kendall tau. In particular the following convex optimisation is used,

$$\begin{array}{ll} \min & \frac{1}{2}\|\mathbf{w}\|^2 + C\sum \varepsilon_{i,j,k} \\ \text{s.t.} & \forall(\mathbf{x}_i, \mathbf{x}_j) \in h_1 : \langle \mathbf{w}, \mathbf{x}_i \rangle \geq \langle \mathbf{w}, \mathbf{x}_j \rangle + 1 - \varepsilon_{i,j,1} \\ & \cdots \\ & \forall(\mathbf{x}_i, \mathbf{x}_j) \in h_n : \langle \mathbf{w}, \mathbf{x}_i \rangle \geq \langle \mathbf{w}, \mathbf{x}_j \rangle + 1 - \varepsilon_{i,j,n} \\ & \forall i, j, k : \varepsilon_{i,j,k} \geq 0, \end{array}$$

in which C is a user-defined variable which trades off the margin size against training error, $\varepsilon_{i,j,k}$ are slack variables and $\mathbf{w}$ is a weight vector. Here, $h_1, \ldots, h_n$ are rankings such that if $(\mathbf{x}_i, \mathbf{x}_j) \in h_i$ then $\mathbf{x}_i$ is ranked higher than $\mathbf{x}_j$. As with SVMs one can use a weight vector which a linear combination of the examples and this allows for the use of kernels for non-linear feature mappings.

3. Metabolomic Data

This exploratory retrospective study was performed on sera obtained from high level or professional cyclists. Samples were stored in a biobank built in the frame of a large longitudinal medical follow-up of cyclists. This biobank was put under the responsibility of the French Federation of Cycling (FFC). The 3 hormones cortisol, IGF-1, and testosterone, were routinely assayed to detect endocrine disruptions as described for cortisol [21]. From the analysis of variations of cortisol, testosterone and IGF-1 concentrations, we have focussed on 3 types of substances suspected to induce endocrine disruptions, which have indirect repercussions on the general metabolism as it is accessible at the serum level by a NMR-based metabolomic observation. More precisely, use of prohibited compounds can be suspected from variation of cortisol, testosterone or IGF-1 levels in serum from what is considered as normal values.

Most of samples were collected in 2001 and 2002 sporting seasons as previously described by Guinot et al. [21]. 655 samples were chosen to constitute this working biobank, which was built from 250 sportsmen submitted to a longitudinal medical follow-up. All samples were obtained anonymously to comply with the French legislation concerning the scientific follow-up of human populations as published previously.

3.1. **Feature Generation and Processing.** For the complete set of 655 NMR spectra the range of concentrations of IGF-1, cortisol and testosterone are shown in Table 3.1. Each spectrum is composed of 950 measurements of resonance at differing frequencies. Most of the samples for testosterone and IGF-1 fall within the normal range. With cortisol we see close numbers of normal and high sample concentrations. Rather that working solely on the raw spectra, we use PCA with 100 eigenvectors to decorrelate features, and also Haar, Daubechies-4 (Db4) and Daubechies-8 (Db8) multilevel wavelets with $J = 10$ levels, resulting in 100, 953, 1012, 1093 features respectively. In order to use a discrete wavelet transform the input feature vector much be of length 2^p where p is an integer, and we pad our features with the border values for them to satisfy this requirement.

	low		normal		high	
	Conc.	Freq.	Conc.	Freq.	Conc.	Freq.
testosterone	0-3	78	3-9	500	9-13	57
cortisol	0-89	62	89-255	303	255-573	268
IGF-1	0-200	71	200-441	472	441-781	88

Table 1. The concentration ranges (ng/ml) for each hormone into the categories: low, normal and high. Also shown is the number of individuals in the sample within the corresponding range.

Since hormone concentrations are known to vary with age, we include the age in years as an additional feature in all feature sets. Any missing values in the label values result in the corresponding example being omitted, however only a maximum of 24 examples contained missing values depending on the label under consideration. There were several age values that were missing and we replaced these with the mean age from the non-missing values.

4. Computational Analysis

Before ranking the spectra, we study the correlation present in the current set of features to confirm the need for feature selection. To do so we compute a correlation matrix $\mathbf{R}$ for each feature set, where $\mathbf{R}_{ij} = \sum_k \mathbf{X}_{ki}\mathbf{X}_{kj} / \sqrt{\sum_k \mathbf{X}_{ki}^2 \sum_k \mathbf{X}_{kj}^2}$ for zero-mean examples. Figure 1 shows the number of pairs in the correlation matrices for each feature set above different thresholds. The key observation is that the raw spectrum contains a large number of features which are closely correlated. There is an initial drop in the beginning of the curve since the spectra are zeroed out for the range corresponding to water. Nearly 10% of the pairs of features have a correlation greater or equal to 0.99. Wavelet features are less correlated than the raw spectrum ones, however both are still significantly more correlated than features which are uniformly randomly generated.

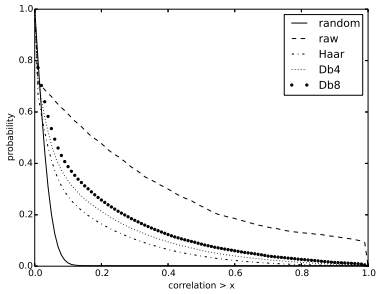

Figure 1. The correlations of the spectra and wavelet coefficients.

Next, PCA is applied to decorrelate the features. The largest 100 eigenvectors of the covariance matrix are computed for the raw and wavelet features, with the covariance matrix defined using $\mathbf{C} = \frac{1}{n}\mathbf{X}^T\mathbf{X}$ for centred examples. We observed that the first 100 eigenvalues model 99.1% of the total variance for the raw spectra and approximately 99.99% of the variance of the wavelets. This implies that the features live predominantly in a low dimensional subspace. Adjacent to this analysis is the reconstruction of the spectra using a filtered subset of size N of the wavelet coefficients. In Figure 2(a) we observe how the error in the filtered spectra vary with N by measuring the Frobenius norm difference between the real and approximated spectra, where the Frobenius norm of $\mathbf{A}$ is $\|\mathbf{A}\|_F^2 = \sum_{ij} \mathbf{A}_{ij}^2$. The initial error is 2139.4 (the norm of the original spectra) and this falls to 372 on average for $N = 100$. Haar wavelets result in the lowest error and the highest are found with the Db8 wavelets. In Figure 2(b) we show an example of the reconstruction of a spectrum after filtering Haar wavelet coefficients. One can see that minor details are removed and large peaks are preserved which is one would expect when performing filtering. We later use an L1 regularised SVM in conjunction with TreeRank Forests for feature selection based on predictive ability, rather than variance as with filtering for example.

4.1. **Learning Setup.** As previously discussed, a useful way of viewing the problems above is as a set of bipartite ranking ones over the class of the corresponding

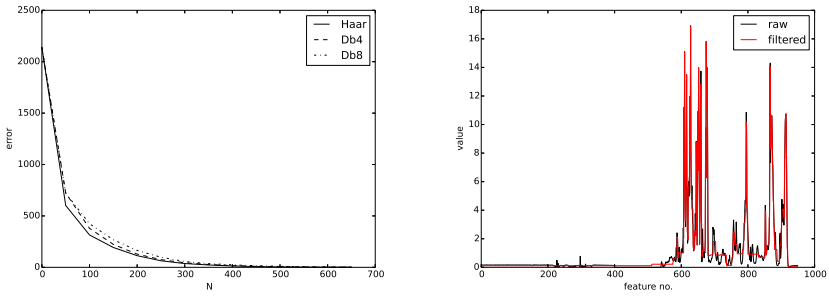

(a) Filter error (b) Reconstructed spectrum

FIGURE 2. An illustration of the errors introduced using wavelet filtering.

hormone (low, normal and high concentrations). The aim is to see which feature sets are more predictive and also to examine the individual features which are useful for learning. We use the following sets of features: raw spectra (denoted `raw`), the PCA transformed data (`pca`), the Daubechies wavelet of length 4 and 8 (denoted by `Db4`, `Db8` respectively), and the Haar wavelet (`Haar`). Each set of features is standardised by centring and scaling each feature to have unit norm. For ranking the following algorithms are used: Ranking Forests (`TRF`), RankBoost (`RB`) and the Ranking SVM (`RSVM`). TreeRank is excluded from the experimentation since in both our preliminary work and in [8] Ranking Forests consistently improves over TreeRank.

4.2. **Bipartite Ranking.** Each approach is evaluated by measuring AUC using 3-fold cross validation and averaging the results. Within each of these folds, model selection is performed using an inner 5-fold cross validation loop. At each cross validation split we stratify the labels to ensure that each split has a distribution of labels similar to the others.

The parameters for the ranking methods are chosen by selecting those with the highest AUC in conjunction with the inner cross validation loop. For RankBoost we choose the number of iterations from $\{10, 50, 100\}$ and the number of weak rankers from $\{5, 10, 20\}$. With the Ranking SVM we use the Radial Basis Function (RBF) kernel defined by $\kappa(\mathbf{x}, \mathbf{x}') = \exp(-\gamma\|\mathbf{x} - \mathbf{x}'\|^2)$. Parameters are chosen as $C \in \{2^0, 2^1, 2^2\}$ and $\gamma \in \{2^{-3}, 2^{-2}, 2^{-1}\}$. For Ranking Forests the minimum split size of a node is set to 50, the maximum depth is fixed at 10 and we use 20 trees. Randomisation improves the performance of TreeRank Forest and so we use bootstrap sampling of the training set on sizes $\{0.5, 0.75, 1.0\}$ of the number of training examples, and similarly we select random feature proportions $\{0.5, 0.75, 1.0\}$. In conjunction with TreeRank Forest we use three LeafRank algorithms: CART, RBF SVM and the L1 penalised SVM denoted `CART`, `RBF` and `L1` in the results. The latter is particularly interesting in terms of finding the relevant features. For each LeafRank algorithm we perform model selection using 3-fold cross validation. For CART we select the maximum depth from $\{2, 4, 8\}$ and fix the minimum split size to 30. In the case of the SVM, $C \in \{2^{-5}, 2^{-3}, \ldots, 2^7\}$ and with the RBF SVM, $\sigma \in \{2^{-5}, 2^{-3}, \ldots, 2\}$.

	IGF-1			cortisol			testosterone		
	low	normal	high	low	normal	high	low	normal	high
	raw								
CART-TRF	.79 (.01)	.72 (.06)	.71 (.01)	.63 (.07)	.68 (.03)	.67 (.04)	.66 (.04)	.57 (.01)	.63 (.03)
L1-TRF	.79 (.02)	.69 (.03)	**.79** (.05)	.71 (.13)	.73 (.01)	**.78** (.03)	**.69** (.07)	**.60** (.03)	.67 (.06)
RBF-TRF	**.80** (.02)	.70 (.01)	.73 (.03)	.71 (.10)	**.74** (.02)	.76 (.04)	.68 (.06)	.57 (.02)	.68 (.04)
RB	.74 (.04)	**.72** (.03)	.77 (.01)	.71 (.08)	.72 (.01)	.76 (.02)	.61 (.07)	.58 (.04)	**.71** (.03)
RSVM	.71 (.03)	.63 (.02)	.71 (.05)	**.74** (.10)	.69 (.01)	.73 (.05)	.66 (.04)	.56 (.02)	.70 (.01)
	pca								
CART-TRF	.68 (.01)	.67 (.04)	.68 (.02)	.65 (.02)	.62 (.02)	.61 (.01)	.62 (.06)	.50 (.04)	.48 (.02)
L1-TRF	**.76** (.01)	.68 (.01)	**.79** (.04)	.69 (.11)	.69 (.03)	.75 (.04)	.63 (.06)	.57 (.02)	.67 (.04)
RBF-TRF	.74 (.03)	.65 (.02)	.78 (.02)	**.70** (.08)	**.72** (.01)	**.77** (.04)	**.66** (.06)	**.61** (.02)	**.70** (.02)
RB	.71 (.02)	**.68** (.05)	.73 (.02)	.54 (.04)	.68 (.03)	.72 (.02)	.60 (.09)	.53 (.03)	.60 (.03)
RSVM	.69 (.01)	.62 (.02)	.74 (.06)	.68 (.08)	.70 (.02)	.76 (.05)	.64 (.05)	.59 (.01)	.67 (.02)

TABLE 2. The AUC values of the ranking methods using the raw and PCA spectra with standard deviations in parentheses and best results in bold.

	IGF-1			cortisol			testosterone		
	low	normal	high	low	normal	high	low	normal	high
	Db4								
CART-TRF	.76 (.03)	.70 (.02)	.71 (.02)	.67 (.09)	.64 (.04)	.70 (.04)	.61 (.04)	.58 (.02)	.61 (.03)
L1-TRF	.74 (.03)	.70 (.01)	.74 (.04)	.69 (.09)	.68 (.02)	.74 (.03)	.66 (.04)	.57 (.01)	.70 (.08)
RBF-TRF	.77 (.03)	**.71** (.01)	**.75** (.06)	.66 (.06)	**.71** (.01)	.76 (.03)	.66 (.02)	**.59** (.02)	.68 (.04)
RB	.73 (.02)	.71 (.01)	.75 (.04)	.65 (.04)	.70 (.01)	.74 (.01)	.63 (.02)	.55 (.03)	.66 (.06)
RSVM	**.78** (.03)	.69 (.01)	.73 (.05)	**.70** (.09)	.70 (.01)	**.77** (.04)	**.69** (.03)	.57 (.02)	**.72** (.04)
	Db8								
CART-TRF	.76 (.02)	**.72** (.02)	.70 (.03)	.64 (.08)	.65 (.01)	.68 (.02)	.61 (.05)	.56 (.05)	.57 (.04)
L1-TRF	.75 (.03)	.72 (.01)	**.76** (.03)	.66 (.08)	.68 (.01)	.74 (.01)	.64 (.05)	**.58** (.03)	.65 (.03)
RBF-TRF	.77 (.03)	.71 (.01)	.74 (.05)	.69 (.10)	.70 (.01)	.76 (.03)	.66 (.05)	.57 (.01)	.67 (.02)
RB	.73 (.03)	.67 (.02)	.73 (.06)	.62 (.08)	.68 (.04)	.69 (.02)	**.68** (.02)	.58 (.02)	.64 (.05)
RSVM	**.78** (.02)	.69 (.01)	.72 (.05)	**.70** (.10)	**.70** (.01)	**.76** (.04)	.68 (.02)	.56 (.02)	**.69** (.03)
	Haar								
CART-TRF	.72 (.05)	**.72** (.01)	.68 (.07)	.62 (.03)	.65 (.02)	.71 (.00)	.61 (.03)	.52 (.03)	.58 (.08)
L1-TRF	.76 (.05)	.70 (.01)	.73 (.06)	.68 (.06)	.69 (.01)	.74 (.03)	.64 (.03)	.58 (.01)	.64 (.09)
RBF-TRF	.77 (.03)	.71 (.01)	**.75** (.06)	.70 (.08)	.70 (.01)	**.77** (.02)	.64 (.01)	**.61** (.02)	.66 (.03)
RB	.77 (.05)	.70 (.03)	.74 (.05)	.60 (.07)	.70 (.04)	.72 (.01)	.60 (.03)	.52 (.02)	.66 (.03)
RSVM	**.78** (.03)	.69 (.01)	.74 (.05)	**.71** (.10)	**.70** (.01)	.76 (.04)	**.67** (.03)	.57 (.01)	**.72** (.04)

TABLE 3. The AUC values of the ranking methods using the wavelet features with standard deviations in parentheses and best results in bold.

Results are presented in Tables 2 and 3. Consider first the best results for each concentration, and surprisingly these are given by the raw spectra. Evidently, the high correlations present in these features do not adversely affect the results. In contrast, moving to PCA gives slightly worse performance in general. We see that high and low concentrations of the hormones are easier to rank than those in the normal range as we might expect since the predicted scoring functions for

the concentrations are bisected more naturally in these cases. IGF-1 is the easiest hormone to rank followed by cortisol and then testosterone. When we look at the ranking methods for the raw spectra, `L1-TRF` gives AUCs closest to the best in every case, perhaps due to its ability to perform feature selection in addition to ranking. Notice that the ranking SVM performs well on the wavelet features, improving over the other methods in most cases.

IGF-1		cortisol		testosterone	
low	high	low	high	low	high
950	950	567	850	235	711
728	777	694	801	321	677
919	722	547	892	887	728
650	808	315	587	645	920
154	554	606	263	566	695
734	305	305	713	732	815
676	247	950	883	304	785
228	698	548	746	757	696
719	645	884	299	746	923
714	732	709	268	772	620

TABLE 4. The 10 most relevant features used by `L1-TRF` on the raw spectra.

Next we attempt to learn about the most important features used for TreeRank Forest based on an approach given in [7]. We focus on `L1-TRF` since it gives strong results using the raw spectra relative to the other ranking methods. To find important features we compute for each ranking tree a marginal gain in AUC given by splitting at an internal node, denoted by $\Delta\text{AUC}(m)$ for a node m. One uses the square of the marginal AUC to weight the weight vector of the SVM and then sums these quantities to give an indication of the overall importance of each feature for each tree. Formally, the vector of feature importances $\hat{\mathbf{w}}$ can be written as

$$\hat{\mathbf{w}} = \sum_{\text{internal nodes } m} (\Delta\text{AUC}(m))^2 \cdot |\mathbf{w}_m|,$$

where $\mathbf{w}_m$ is the weight vector for node m and $|\cdot|$ is the element-wise absolute value of a vector. For a forest of trees, one simply sums the weight vector for each tree and then normalises so that the resulting vector has unit norm. The resulting zero-based indices are presented in Table 4 for high and low hormone concentrations, since these resulted in the highest AUC scores in Table 3. Of note here is that there is little overlap in terms of the features used for the same hormone aside from the feature corresponding to age, 950, in the case of IGF-1. Age also appears to useful for ranking concentrations of cortisol. When we look at the first 100 features then there are 35, 31 and 24 in common for low and high concentrations of IGF-1, cortisol and testosterone respectively. This appears to point to the relevance of using difference ranking schemes to detect abnormal concentrations. We leave a more medically centred interpretation for informed readers.

5. Conclusions

We explored the scoring problem in metabolomics in the context of predicting elevated metabolite presence in blood serum analysed using NMR spectra. Such a

problem was more informatively split into multiple bipartite ranking problems as compared to the traditional classification or regression approach. On a cohort of 655 NMR spectra of French professional or near professional cyclists, the concentration ranges of testosterone, cortisol and IGF-1 were predicted using bipartite ranking methods in an extensive computational analysis. We showed that the raw spectra were more predictive in conjunction with the ranking methods relative to those generated using PCA and wavelet features. Finally we used a novel method to study feature important and highlighted the most important features used for ranking with the L1-penalised SVM TreeRank Forest.

Acknowledgements

The authors are grateful to Alain Paris of AgroParisTech for granting access to the metabolomics database.

References

[1] C. Abraham, G. Biau, and B. Cadre. On the kernel rule for function classification. *Annals of the Institute of Statistical Mathematics*, 58(3):619–633, 2006.

[2] T. Alexandrov, J. Decker, B. Mertens, A. Deelder, R. Tollenaar, P. Maass, and H. Thiele. Biomarker discovery in MALDI-TOF serum protein profiles using discrete wavelet transformation. *Bioinformatics*, 25(5):643, 2009.

[3] R. N. Bracewell and R. Bracewell. *The Fourier transform and its applications*, volume 31999. McGraw-Hill New York, 1986.

[4] L. Breiman, J. Friedman, R. Olshen, and C. Stone. *Classification and regression trees*. Wadsworth and Brooks, 1984.

[5] K. Bryan, P. Cunningham, and L. Brennan. Using Random Forest for Classification and Feature Selection for Metabolomics Data. 2007.

[6] S. Clémençon, M. Depecker, and N. Vayatis. Nonparametric scoring methods as a support decision tool for medical diagnosis The TreeRank algorithm and its variants. In *ECML PKDD Workshop on Knowledge Discovery in Health Care and Medicine*, 2011.

[7] S. Clémençon, M. Depecker, and N. Vayatis. Adaptive partitioning schemes for bipartite ranking. *Machine Learning*, 83(1):31–69, 2011.

[8] S. Clémençon, M. Depecker, and N. Vayatis. Ranking forests. *Journal of Machine Learning Research*, 14:39–73, 2013.

[9] A. d'Aspremont, L. El Ghaoui, M. I. Jordan, and G. R. Lanckriet. A Direct Formulation for Sparse PCA Using Semidefinite Programming. Technical Report UCB/CSD-04-1330, EECS Department, University of California, Berkeley, USA, Cambridge, MA, Jun 2004.

[10] I. Daubechies. Ten lectures on wavelets (CBMS-NSF regional conference series in applied mathematics). 1992.

[11] R. Davis, A. Charlton, J. Godward, S. Jones, M. Harrison, and J. Wilson. Adaptive binning: An improved binning method for metabolomics data using the undecimated wavelet transform. *Chemometrics and intelligent laboratory systems*, 85(1):144–154, 2007.

[12] C. Dhanjal, S. Gunn, and J. Shawe-Taylor. Efficient sparse kernel feature extraction based on partial least squares. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 31(8):1347–1361, 2009.

[13] C. Domange, C. Canlet, A. Traoré, G. Biélicki, C. Keller, A. Paris, and N. Priymenko. Orthologous Metabonomic Qualification of a Rodent Model Combined with Magnetic Resonance Imaging for an Integrated Evaluation of the Toxicity of Hypochœris radicata. *Chemical Research in Toxicology*, 21(11):2082–2096, 2008.

[14] T. Ebbels, J. Lindon, and M. Coen. Processing and modeling of nuclear magnetic resonance (NMR) metabolic profiles. *Methods in molecular biology (Clifton, NJ)*, 708:365–388, 2011.

[15] M. Eliasson, S. Rannar, and J. Trygg. From Data Processing to Multivariate Validation-Essential Steps in Extracting Interpretable Information from Metabolomics Data. *Current Pharmaceutical Biotechnology*, 12(7):996–1004, 2011.

[16] A. Fernie, R. Trethewey, A. Krotzky, and L. Willmitzer. Metabolite profiling: from diagnostics to systems biology. *Nature Reviews Molecular Cell Biology*, 5(9):763–769, 2004.

[17] F. Ferraty and P. Vieu. *Nonparametric functional data analysis: theory and practice*. Springer Verlag, 2006.
[18] Y. Freund, R. Iyer, R. E. Schapire, and Y. Singer. An efficient boosting algorithm for combining preferences. *The Journal of machine learning research*, 4:933–969, 2003.
[19] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. In P. Vitányi, editor, *Computational Learning Theory*, volume 904 of *Lecture Notes in Computer Science*, pages 23–37. Springer Berlin / Heidelberg, London, UK, 1995.
[20] J. Friedman, T. Hastie, and R. Tibshirani. Special invited paper. additive logistic regression: A statistical view of boosting. *Annals of statistics*, pages 337–374, 2000.
[21] M. Guinot, M. Duclos, N. Idres, J. Souberbielle, A. Megret, and Y. Le Bouc. Value of basal serum cortisol to detect corticosteroid-induced adrenal insufficiency in elite cyclists. *European journal of applied physiology*, 99(3):205–216, 2007.
[22] I. M. Guyon, S. R. Gunn, M. Nikravesh, and L. Zadeh, editors. *Feature Extraction, Foundations and Applications*. Springer, 2006.
[23] T. Hastie, A. Buja, and R. Tibshirani. Penalized discriminant analysis. *The Annals of Statistics*, pages 73–102, 1995.
[24] T. Hastie, R. Tibshirani, J. Friedman, and J. Franklin. The elements of statistical learning: data mining, inference and prediction. *The Mathematical Intelligencer*, 27(2):83–85, 2005.
[25] L. Hoegaerts, J. A. K. Suykens, J. Vandewalle, and B. D. Moor. Primal space sparse kernel partial least squares regression for large scale problems. In *Proceedings of the IEEE International Joint Conference on Neural Networks*, volume 1, pages 561–566, 2004.
[26] E. Holmes, I. Wilson, and J. Nicholson. Metabolic phenotyping in health and disease. *Cell*, 134(5):714–717, 2008.
[27] H. Hotelling. Analysis of a complex of statistical variables into principle components. *Journal of Educational Psychology*, 24:417–441 and 498–520, 1933.
[28] T. Joachims. Optimizing search engines using clickthrough data. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 133–142. ACM, 2002.
[29] S. Mahadevan, S. L. Shah, T. J. Marrie, and C. M. Slupsky. Analysis of Metabolomic Data Using Support Vector Machines. *Analytical Chemistry*, 80(19):7562–7570, 2008. PMID: 18767870.
[30] S. Mallat. *A wavelet tour of signal processing*. Academic Pr, 1999.
[31] B. Moghaddam, Y. Weiss, and S. Avidan. Spectral bounds for sparse pca: Exact and greedy algorithms. In *Advances in neural information processing systems*, pages 915–922, 2005.
[32] M. Momma and K. P. Bennett. Sparse Kernel Partial Least Squares Regression. In B. Schölkopf and M. K. Warmuth, editors, *Proceedings of the 16th Annual Conference on Computational Learning Theory*, pages 216–230, 2003.
[33] J. Nicholson, J. Lindon, E. Holmes, et al. Metabonomics': understanding the metabolic responses of living systems to pathophysiological stimuli via multivariate statistical analysis of biological NMR spectroscopic data. *Xenobiotica*, 29(11):1181–1189, 1999.
[34] J. Ramsay. Functional data analysis. *Encyclopedia of Statistical Sciences*, 1997.
[35] F. Rohart, N. Villa-Vialaneix, A. Paris, C. Canlet, J. Molina, D. Milan, B. Laurent, and M. SanCristobal. Phenotypic prediction based on metabolomic data : lasso vs Bolasso, primary data vs wavelet data. In *In Proceedings of WCGALP 2010*, 2010.
[36] A. Roux, D. Lison, C. Junot, and J. Heilier. Applications of liquid chromatography coupled to mass spectrometry-based metabolomics in clinical chemistry and toxicology: A review. *Clinical biochemistry*, 44:119–135, 2010.
[37] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.
[38] J. Trygg, E. Holmes, and T. Lundstedt. Chemometrics in metabonomics. *Journal of proteome research*, 6(2):469–479, 2007.
[39] H. Wold. Estimation of principal components and related models by iterative least squares. *Multivariate analysis*, pages 391–420, 1966.
[40] J. Xia, X. Wu, and Y. Yuan. Integration of wavelet transform with PCA and ANN for metabolomics data-mining. *Metabolomics*, 3(4):531–537, 2007.
[41] J. Zhu and T. Hastie. Kernel logistic regression and the import vector machine. *Advances in Neural Information Processing Systems*, 2:1081–1090, 2002.

[42] H. Zou, T. Hastie, and R. Tibshirani. Sparse principal component analysis. *Journal of computational and graphical statistics*, 15(2):265–286, 2006.